

Claims

We claim:

1 1. A method for configuring and rebuilding a redundant array of independent disks,
2 comprising:

3 configuring the array with D disk drives of B physical blocks each;

4 allocating N user data and redundant data blocks to each disk drive; and

5 allocating F free blocks as hot spare space to each disk drive, where $N + F$
6 $\leq B$, and $((D - M) \times F) \geq N$ to enable rebuilding of data and redundant blocks of
7 a failed disk drive in the free blocks of the remaining disk drives after M disk drive
8 failures.

1 2. The method of claim 1 further comprising:

2 detecting M disk drive failures;

3 generating new user data and redundant data to recover from the M disk
4 failures;

5 moving the new data and redundant data into the hot spare space; and

6 rearranging all of the data and redundant blocks of the D-M disk drives to
7 rebuild the array with fewer disks and an identical level of redundancy.
8

1 3. A method of claim 1 further comprising:

2 generating new user data and redundant to recover from a single failed disk
3 drive;

4 moving the new data and redundant data into a part of the hot spare space;

5 and

rearranging all of the data and redundant blocks of the remaining disk drives to rebuild the array with fewer disks and an identical level of redundancy so that remaining hot spare space which can accommodate M-1 additional disk drive failures.

4. A method of claim 1 further comprising:

generating new data and redundant data for concurrent disk drive failures of more than one disk drive;

moving the new data and redundant data into a part of the hot spare space; and

rearranging all of the data and redundant blocks of the remaining disk drives to rebuild the array with fewer disks and an identical level of redundancy.

5. The method of claim 1 further wherein where M is equal to one, and further comprising:

generating new data and redundant data for a failure of exactly one disk drive;

moving the new data and redundant data into a part of the hot spare space; and

rearranging all of the data and redundant blocks of the remaining disk drives to rebuild the array with fewer disks and an identical level of redundancy.

6. The method of claim 1 wherein the redundant data are mirror blocks.

7. The method of claim 1 wherein the redundant data are parity blocks.

1 8. The method of claim 1 wherein all disk drives except failed disk drives are
2 actively used while recovering from the M disk drive failures.

1 9. The method of claim 2 wherein the generating uses an exclusive OR operation.

1 10. A redundant array of independent disks with hot spare space, comprising:
2 D disk drives of B physical blocks each;
3 N user data and redundant data blocks allocated to each of the D disk drives;
4 F free blocks allocated as hot spare space to each of the D disk drive, where
5 $N + F \leq B$, and $((D - M) \times F) \geq N$; and
6 means for rebuilding the data and redundant blocks of a failed disk drive in
7 the free blocks of the remaining disk drives after M disk drive failures.